

09/787198

2 Rec'd PCT/PTO 15 MAR 2001

1

PROCESSING PLATFORM

The present invention relates to a processing platform and particularly to telecommunications service platform using a number of loosely coupled  
5 subsystems.

Platforms used, for example, to implement advanced services in a telecommunications network having an IN (intelligent network) architecture, may comprise a number of loosely coupled subsystems linked by a high speed network. This structure enhances the capacity and the resilience of the platform. The  
10 resources of the platform are managed by a centralised management system that is responsible for allocating particular resources to given call and for reconfiguring the system, for example, in the event of one of the system components failing. As conventionally implemented, such a platform suffers the disadvantages of having a high management overhead and having poor scalability, that is, the platform is not  
15 readily adaptable to provide increased capacity.

According to a first aspect of the present invention, there is provided a communications service platform comprising a multiplicity of loosely coupled subsystems, each of the subsystems including respective service processing resources and a respective resource locator, each resource locator including  
20 means for communicating to others of the resource locators its identity and the availability of the resources in the respective subsystem, and means for receiving identity data and resource availability data for resource locators in others of the subsystems.

This aspect of the invention provides a service platform which maintains  
25 the resilience and capacity of a distributed processing architecture while reducing the management overheads and providing greatly improved scalability. This is achieved by providing each subsystem with a resource locator which advertises the resources of the relevant subsystem and also listens to information from other subsystem. In this way the task of resource management and allocation is  
30 distributed between the subsystems, and the system as a whole is provided with a mechanism for recognising and responding to the addition of new subsystems.

The resource locators may be arranged to communicate directly by peer-to-peer signalling. Preferably, however, the platform includes a resource broker and communication between the resource locators is mediated by the resource

broker. The resource broker may be located in one of the service processing subsystems, or may be located in a dedicated platform. The resource broker may include a data interface arranged to receive capability data and interface data from respective resource locators, and a registry arranged to store the said capability  
5 data and interface data. Preferably a resource locator in a subsystem is arranged initially to read capability data and interface data for another subsystem from the resource broker, and subsequently communicates further data directly with the other subsystem using the interface of the subsystem identified in the said interface data.

10 The system developed by the present inventors is not limited to use in communications systems but may also be used in a general computing environment.

According to a second aspect of the present invention there is provided a computing platform comprising a multiplicity of loosely coupled computing  
15 subsystems, each of the said subsystems including respective data processing resources and a respective resource locator arranged to advertise its identity and the loading of the respective resources and to receive resource signalling from others of the resource locators .

According to a third aspect of the present invention, there is provided a  
20 method of operating a communications system, the system including a multiplicity of processing subsystems and a network interconnecting the multiplicity of subsystems, the method comprising;

a) communicating from a resource locator in a respective one of the multiplicity of subsystems to resource locators in others of the multiplicity of  
25 subsystems data indicating the identity of the said one subsystem and the availability of resources in the said one subsystem

b) repeating step (a) for each other of the multiplicity of subsystems:

c) when one of the multiplicity of subsystems, in the course of processing a call, requires resources not present locally in the said subsystem:

30 i) identifying from the said data communicated to the resource locator of the said one subsystem another subsystem having the said resources;  
ii) accessing the said subsystem via the network.

According to a fourth aspect of the present invention, there is provided a communications system comprising:

- a plurality of call processing subsystems;
- a network interconnecting the plurality of call processing subsystems;
- a resource broker connected to the network, the resource broker including
  - a data interface arranged to receive capability
- 5 data and interface data from respective call processing subsystems, and
  - a registry arranged to store the said capability data and interface data.

The term "call processing subsystem" is used here broadly to encompass systems ancillary to the processing of a call, such as, e.g., an Email server, a mobility application platform or an interactive voice response (IVR) platform, as well as systems, such as a communications server, which are directly involved in the handling of a call.

- 15 Systems embodying the present invention will now be described in further detail, by way of example only, with reference to the accompanying drawings in which;

Figure 1 is a schematic of network embodying the invention;

Figure 2 is a diagram showing the structure of a service control point in

- 20 Figure 1;

Figure 3 is a schematic of a network employing peer-to-peer signalling;

Figure 4 is a schematic of a network employing a broker;

Figure 5 is a diagram showing interactions between a broker and components;

- 25 Figure 6 is a schematic showing an architecture implementing the invention;

Figure 7 is a class diagram for an implementation of the architecture of Figure 6;

Figure 8 is a message flow diagram showing a first scenario;

- 30 Figure 9 is a message flow diagram showing a second scenario;

Figure 10 is a schematic of an implementation of the architecture of Figure 6 across different networks;

Figure 11 is a diagram illustrating different interface types in a system embodying the invention;

Figure 12 is a diagram showing the structure of a resource broker;

Figure 13 shows a system in which components incorporate both type 1.0 and type 2.0 interfaces

Figure 14 shows a personal mobility service implemented using the architecture of Figure 6.

A telecommunications network which uses an IN (Intelligent Network) architecture includes a service control point 1, which is also termed herein the Network Intelligence Platform (NIP). The service control point 1 is connected to trunk digital main switching units (DMSU's) 2,3 and to digital local exchanges (DLE's) 4,5. Both the DMSU's and the DLE's function as service switching points (SSP's). At certain points during the progress of a call, the SSP's transfer control of the call to the service control point. The service control point 1 carries out functions such as number translation and provides a gateway to additional resources such as a voice messaging platform. The network also includes a second service control point 6, which in this example runs customer mobility applications. The two service control points are connected to each other and to the other components via a broadband signalling network 7.

Figure 2 shows the structure of the service control point 1. A service management server is connected via an FDDI optical fibre LAN 21 to an overload control server (OCS) and to transaction servers (TS). The overload control server monitors call rates and initiates action to protect the SCP and other elements of the network from overload. For example the OCS may send a call-gapping instruction to an originating exchange. Additionally or alternatively, a call-gapping instruction may be sent to the broker, so that the broker rations the availability, e.g., of switching resources at the originating exchange so as to protect other resources downstream of the originating exchange. The transaction servers implement advanced service control functions such as, for example number translation and call plans. The OCS and transaction servers are connected via a second FDDI LAN 22 to communications servers (CS) which are connected to an SS7 (ITU Signalling System no. 7) signalling network. A global data server (GDS) is also connected to this second LAN. The GDS records call statistics and may also be used in the operation of overload control functions.

Figure 3 is a simplified schematic of the network of Figure 1, illustrating a first implementation of the present invention. The Figure shows a first SCP 31, a second SCP 32 and a DMSU 33. These components constitute loosely coupled subsystems of a platform which runs one or more IN services. The subsystems are  
5 connected by a broadband signalling network 34, for example an FDDI LAN. As indicated by the key to the Figure, each subsystem includes call processing resources of different types and also a resource locator 35. The first SCP includes resource types 1 and 2 which might correspond, for example, to a VPN (virtual private network) control function and to a number translation function  
10 respectively. The second SCP includes type 1 resources, and the DMSU includes resource type 3, for example a call switching function. Each locator advertises to all other locators, the identity and location of the subsystem, the resources available at that location and, for example, the loading of the resources. For example the locator in the first SCP 31, when the SCP is initiated, broadcasts to  
15 the other resource locators, a message in the form (SCP1; address1; type 1, 50%; type 2, 90%) where the first field is the identity, the second field is the address, and the subsequent fields are resource types and include a measure of resource loading. This message is received and stored by the resource locators in the other subsystems, and they in turn broadcast resource messages which are received by  
20 the resource locator in the SCP 31. The step of broadcasting resource messages is repeated periodically, with a frequency chosen to balance the need of the system to re-balance resources after a subsystem failure, or after incremental changes in the proportion of free resources in a subsystem, against the need to minimise signalling overheads. In the present example, resource locators re-  
25 transmit a resource message every 10 seconds. Then, for example, if a call originating at the DMSU is making use of VPN resources in the first SCP, and the first SCP fails, the failure will become apparent within, at most 10 seconds, and the DMSU resource locator may identify alternative resources in the second SCP. The absence of an expected update is interpreted by the resource locators as an  
30 indication that the corresponding resources have become unavailable. In addition, subsystems may update the resource locators in response to events within that subsystem. For example, a subsystem may be programmed to update the resource locators whenever its loading changes by more than a predetermined threshold, e.g. 20%.

Figure 4 is a schematic illustrating a second implementation of the invention. In this example, in addition to the first and second SCP's and the DMSU described above with reference to Figure 3, the system also includes a resource broker 40. The resource broker includes, in addition to its own resource locator 41, a data store 42 which holds a registry of data communicated by other subsystems, and an authentication processor 43 which authenticates data received from the subsystems prior to entering the data in the registry. The broker may be implemented on a commercially available system such as that available from Visigenics as the Visigenics ORB (object request broker) running, for example, on a Digital SPARC Ultra (trademark) workstation. As shown in Figure 5, a subsystem, here termed a "component", must first authenticate itself with the broker, for example using a password or a digital signature, and agree a security mechanism for further exchanges. The component then registers its interfaces with the broker. For example, a communications server might register an INAP communications channel at a specified network address. From this point, the component can discover the interfaces of other components within the system by reference to the registry in the broker. Once the component understands the interfaces of other components it can communicate freely with the other components without further reference to the broker.

Figure 12 shows in further detail the structure of the resource broker 40. It comprises a resource registry, a service registry, and modules for authentication, registration and discovery. Tables 1 and 2 below show examples of the contents of the resource registry and service registry respectively. Each resource in the resource registry may have linked with it a number of entries in the service registry.

TABLE 1

Name: Athena
Processor: Digital SPARC Ultra
Address: 132.14.32.71
Function: Call Control Server

TABLE 2

Call Control:INAP
-------------------

no. RANGE 64XXXX-67XXXX
Network Operator: BT
Capacity: 100 cps
Security parameters:
authentication level
signature
Interface: VIPER 1.0

Figure 6 shows an overview of an architecture that is implemented using a resource broker as described above with reference to Figures 4 and 5. The broker and the network used by the subsystems to communicate with the broker together provide a brokered environment 60. In this example, the subsystems connected to the brokered environment 60 include a PSTN network 61, a gatekeeper 62 for voice over IP (internet protocol) services, fax resources 63, an Email server 64 and a voice mail server 65. A number of applications use the resources of the subsystems 61-65. These applications locate the appropriate resources using the resource broker contained within the brokered environment 60. This implementation is termed by the inventors the "VIPER" architecture. The architecture provides for two types of interfaces, termed by the inventors VIPER 1.0 and VIPER 2.0. As illustrated in Figure 11, VIPER 2.0 interfaces provide for communication between subsystems via an object bus 110. Subsystems using VIPER 2 interfaces request resources from a resource broker on a per call basis and communicate using, e.g., the CORBA protocol. VIPER 1 interfaces bypass the object bus and use specific protocols such as INAP to communicate with other systems. Subsystems with VIPER 1 interfaces register and discover resources and interface details with the resource broker when the subsystem is initialised, but do not communicate with the resource broker for subsequent calls. Such subsystems using VIPER 1 interfaces then communicate with other subsystem using protocols specific to the subsystem, for example INAP in the case of a communications server communicating with a transaction server.

In operation, a subsystem such as a communications server may initially download copies of service and resource registry data from the resource broker to form a locally cached copy of the resource broker. For example, in Figure 11, the

communications server CS may optionally include a locally cached broker (LCB), as shown in dashed lines in the Figure. Then discovery operations may be carried out using the locally cached copy of the resource broker. In this case, even where a VIPER 2.0 interface is used, it is not necessary for signalling to pass between the

5 communications server and the remote broker on every call. Instead, data passes between the remote broker and the communications server only intermittently when it is necessary to refresh the locally cached resource broker. Although the communications server, if it employs a VIPER 2.0 interface, still interrogates the resource broker for each new call, it is in general the locally cached broker that is

10 used for this purpose.

Figure 7 is a class diagram showing an implementation of the architecture of Figure 6. This diagram is generated using the Rational ROSE software engineering tool and provides a basis, using that tool, for generating, e.g., Java code implementing the invention. Rational ROSE is available commercially from

15 Rational Software Corporation. In this implementation, the subsystems that interact using the broker are termed "plugins".

Figure 8 shows a network operator bringing a VIPER broker, a VIPER Cambridge Transaction Server and a VIPER Cambridge Communications Server into service. "Cambridge" is the name given to the SCP illustrated in Figure 2. After

20 registration and discovery is complete, an incoming call triggers an INAP (Intelligent Network Application Protocol) IDP (initial detection point) at a DLE. The DLE is referenced GPT\_SSP in the diagram. The call is cut-through the VIPER middleware. Since the communications server CS and transaction server TS have both registered VIPER 1.0 interfaces with the resource broker, the CS does not

25 have to ask the broker each time a call is received for the address of a suitable service resource or "plugin". Instead the CS and TS communicate directly using an INAP interface and bypassing the object bus. After location lookup is performed, an INAP connect is also cut-through the VIPER middleware back to the SSP. The following events and messages are shown in Figure 8:

30

1: Network operator brings the VIPER broker into service.

2: VIPER broker registers its services with itself if necessary.



3: Network operator brings the VIPER Cambridge Transaction Server into service.

4: VIPER Cambridge Transaction Server registers itself as a PlugIN with the VIPER broker.

5

5: VIPER Cambridge Transaction Server registers the services it supports with the VIPER broker.

6: Network operator brings the VIPER Cambridge Communications Server into  
10 service.

7: VIPER Cambridge Communications Server registers itself as a PlugIN with the VIPER broker.

15 8: VIPER Cambridge Communications Server registers the services it supports with the VIPER broker.

9. VIPER Cambridge Transaction Server requests the address of the PlugIN associated with 'name' (in this case, VIPER Cambridge Communications Server).  
20

10. VIPER Cambridge Communications Server requests the address of the PlugIN associated with 'name' (in this case, VIPER Cambridge Transaction Server).

11. GPT SSP sends an INAP Initial Detection Point (IDP) to the VIPER Cambridge  
25 Communications Server.

12. VIPER Cambridge Communications Server sends a INAP IDP to the VIPER Cambridge Transaction Server.

30 13. VIPER Cambridge Transaction Server does an internal location lookup, then sends an INAP connect to the VIPER Cambridge Communications Server.

14. VIPER Cambridge Communications Server sends an INAP connect to the GPT SSP.

Figure 9 shows the message flows involved in a second scenario, with a network operator bringing a VIPER broker, a VIPER Mobility Server and a VIPER Cambridge Communications Server into service. The VIPER mobility server may be, for example, the SCP 6 running a mobility application illustrated in Figure 1. Once registration is complete, an incoming call triggers an INAP IDP into the VIPER middleware. A communications server CS creates a call model (labelled "INAP call" in Figure 9) and passes control of the service to that call model. The call model communicates with the resource broker to identify an application that is capable and ready to provide the resources required by the service. That application, if not already running, is created. The call model then request the application, in relation to this call, to perform a location look-up for the called party. On completion of the look-up, a mobile address returned by the look-up is passed to the communications server and the VIPER middleware send an INAP connect to the SSP.

1: Network operator brings the VIPER broker into service.

2: VIPER broker registers its services with itself if necessary.

3: Network operator brings the VIPER Mobility Server into service.

4: VIPER Mobility Server registers itself as a PlugIN with the VIPER broker.

5: VIPER Mobility Server registers the services it supports with the VIPER broker.

6: Network operator brings the VIPER Cambridge Communications Server into service.

7: VIPER Cambridge Communications Server registers itself as a PlugIN with the VIPER broker.

8: VIPER Cambridge Communications Server registers the services it supports with the VIPER broker.

9. GPT SSP sends an INAP Initial Detection Point (IDP) to the VIPER Cambridge Communications Server.

5 10 & 11. The VIPER Cambridge Communications Server creates a new call model to handle this service and then initiates the Call Model's constructor.

12. The Call Model requests from the VIPER broker the address of the PlugIN (in this case, the VIPER Mobility Server) capable of providing the service described  
10 within the service descriptor.

13, 14 & 15. The Call Model requests from the VIPER Mobility Server the address of the application capable of servicing the request. This causes the VIPER Mobility Server to create a Roaming Application and then initiate the Roaming Application's  
15 constructor. Flow 14 is optional depending on whether the Roaming Application is active or not.

16. The Call Model sends a VIPER 2.0 equivalent of an INAP IDP to the Roaming Application.  
20

17. The Roaming Application performs a location lookup.

18. The Roaming Application sends a VIPER 2.0 Connect to the Call Model.

25 19. The Call Model sends an internal PlugIN connect to the VIPER Cambridge Communications Server.

20. The VIPER Cambridge Communications Server send an INAP connect to the GPT SSP.  
30

Figure 10 illustrates an implementation of the VIPER architecture which provides a service platform which is distributed over different networks that may be located in different continents. In this example a customer at a customer

terminal 100 is registered with a mobility application resident on a first application server connected to a first broadband signalling network e.g. in the UK. The customer terminal is connected via a digital local exchange DLE to a second broadband signalling network, e.g., in the US. The first and second national  
5 networks are interconnected via a jointly operated international network 103. Resource brokers (referenced B) are connected to the networks and the resources available on each network are registered with the resource brokers. The resource brokers communicate data to each other via a global broker GB, so that, for example, the broker B connected to network 102 has in its registry details of  
10 resources on both network 101 and network 102. The communication between brokers may be implemented, e.g., using CORBA or IIOP over IP.

In operation, the customer at terminal 100 registers their current location with the mobility application on the first application server. To do this, they dial the number associated with the service. This triggers an IDP at the SSP. The SSP  
15 requests from a local broker B the address of the application required to service the call. The local broker B communicates with the global broker GB connected to network 103 to obtain interface and address data to allow the SCP to communicate with application server 1. The mobility application may, for example, require the use of an interactive voice response (IVR) system to accept data from  
20 the customer. This resource is provided within network 101 by a first intelligent peripheral IP1. However, the second network includes its own IVR resources provided by a second intelligent peripheral IP2. The mobility application requests discovery of IVR resources from the broker B. After communication with the global broker GB, the broker returns detail of the IP2 that is local to the customer  
25 at the current customer location. The application server 1 uses this information to return an INAP connect message to the SSP to set up a connection between the customer terminal and IP2.